Computational Difficulty of Computing the Density of States

Brielin Brown,[1,2] Steven T. Flammia,[2] and Norbert Schuch[3]

[1] *University of Virginia, Departments of Physics and Computer Science, Charlottesville, Virginia 22904, USA*
[2] *Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada*
[3] *California Institute of Technology, Institute for Quantum Information, MC 305-16, Pasadena, California 91125, USA*

We study the computational difficulty of computing the ground state degeneracy and the density of states for local Hamiltonians. We show that the difficulty of both problems is exactly captured by a class which we call #BQP, which is the counting version of the quantum complexity class QMA. We show that #BQP is not harder than its classical counting counterpart #P, which in turn implies that computing the ground state degeneracy or the density of states for classical Hamiltonians is just as hard as it is for quantum Hamiltonians.

Understanding the physical properties of correlated quantum many-body systems is a problem of central importance in condensed matter physics. The density of states, defined as the number of energy eigenstates per energy interval, plays a particularly crucial role in this endeavor. It is a key ingredient when deriving many thermodynamic properties from microscopic models, including specific heat capacity, thermal conductivity, band structure, and (near the Fermi energy) most electronic properties of metals. Computing the density of states can be a daunting task however, as it in principle involves diagonalizing a Hamiltonian acting on an exponentially large space, though other more efficient approaches which might take advantage of the structure of a given problem are not a priori ruled out.

In this Letter, we precisely quantify the difficulty of computing the density of states by using the powerful tools of quantum complexity theory. Quantum complexity aims at generalizing the well-established field of classical complexity theory to assess the difficulty of tasks related to quantum mechanical problems, concerning both the classical difficulty of simulating quantum systems as well as the fundamental limits to the power of quantum computers. In particular, quantum complexity theory has managed to explain the difficulty of computing ground state properties of quantum spin systems in various settings, such as two-dimensional (2D) lattices [1] and even one-dimensional (1D) chains [2], as well as fermionic systems [3].

We will determine the computational difficulty of two problems: First, computing the density of states of a local Hamiltonian, and second, counting the ground state degeneracy of a local gapped Hamiltonian; in both cases, the result holds even if the Hamiltonian is restricted to act on a 2D lattice of qubits, or on a 1D chain. To this end, we will introduce the quantum counting class #BQP (sharp BQP), which constitutes the natural counting version of the class QMA (Quantum Merlin Arthur) which itself captures the difficulty of computing the ground state energy of a local Hamiltonian [4, 5]. Vaguely speaking, #BQP counts the number of possible "quantum solutions" to a quantum problem that can be verified using a quantum computer. We show that both problems, computing the density of states and counting the ground state degeneracy, are complete problems for the class #BQP, i.e., they are among the hardest problems in this class.

Having quantified the difficulty of computing the density of states and counting the number of ground states, we proceed to relate #BQP to known classical counting complexity classes, and show that #BQP equals #P (under weakly parsimonious reductions). Here, the complexity class #P counts the number of satisfying assignments to any efficiently computable boolean function. This can be a very hard problem which is believed to take exponential time; in particular, it is at least as hard as deciding whether the function has at least one satisfying input, i.e., the complexity class NP. Examples for #P-complete problems (i.e., the hardest problems in that class) include counting the number of colorings of a graph, or computing the permanent of a matrix with binary entries. Phrased in terms of Hamiltonians, what we show is that computing the density of states and counting the ground state degeneracy of a classical spin system is just as hard as solving the same problem for a quantum Hamiltonian.

Quantum complexity classes.—Let us start by introducing the relevant complexity classes. The central role in the following is taken by the *verifier* V, which verifies "quantum solutions" (also called *proofs*) to a given problem. More formally, a verifier checking an n-qubit quantum proof (that is, a quantum state $|\psi\rangle$) consists of a $T = \mathrm{poly}(n)$ length quantum circuit $U = U_T \cdots U_1$ (with local gates U_t) acting on $m = \mathrm{poly}(n)$ qubits, which takes the n-qubit quantum state $|\psi\rangle_I$ as an input, together with $m - n$ initialized ancillas, $|\mathbf{0}\rangle_A \equiv |0 \cdots 0\rangle_A$, applies U, and finally measures the first qubit in the $\{|0\rangle_1, |1\rangle_1\}$ basis to return 1 ("proof accepted") or 0 ("proof rejected"). Then, the class QMA contains all problems of the form: "Decide whether there exists a $|\psi\rangle$ such that $p_{\mathrm{acc}}(V(\psi)) > a$, or whether $p_{\mathrm{acc}}(V(\psi)) < b$ for all $|\psi\rangle$, for some chosen $a - b > 1/\mathrm{poly}(n)$, given that one is the case". Here, the acceptance probability of a state $|\psi\rangle$ is $p_{\mathrm{acc}}(V(\psi)) := \langle\psi|\Omega|\psi\rangle$, with

$$\Omega = (\mathbb{1}_I \otimes \langle\mathbf{0}|_A) U^\dagger (|1\rangle\langle 1|_1 \otimes \mathbb{1}) U (\mathbb{1}_I \otimes |\mathbf{0}\rangle_A)\,, \qquad (1)$$

which we illustrate in Fig. 1.

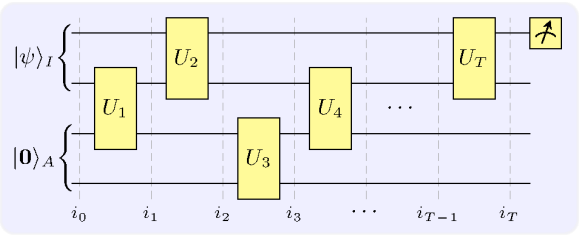

FIG. 1: A QMA verifier consists of a sequence of T local unitary gates acting on the "quantum proof" $|\psi\rangle$ and an ancillary register initialized to $|\mathbf{0}\rangle$. The final measurement on the first qubit returns $|1\rangle$ or $|0\rangle$ to accept or reject the proof, respectively. Transition probabilities can be computed by doing a "path integral" over all intermediate configurations $(i_k)_k$.

The idea behind this definition is that QMA quantifies the difficulty of computing the ground state energy $E_0(H)$ of a local Hamiltonian H up to $1/\mathrm{poly}(n)$ accuracy. Let the verifier be a circuit estimating $\langle\psi|H|\psi\rangle$; then a black box solving QMA problems can be used to compute $E_0(H)$ up to $1/\mathrm{poly}(n)$ accuracy by binary search using a single QMA query. Note also that QMA is the quantum version of the class NP, where one is given an efficiently computable boolean function $f(x) \in \{0,1\}$ and one must figure out if there is an x such that $f(x) = 1$.

The class NP has a natural counting version, known as #P. Here, the task is to determine the *number* rather than the existence of satisfying inputs, i.e., to compute $|\{x : f(x) = 1\}|$. We will now analogously define #BQP, the counting version of QMA. Consider the verifying map Ω of Eq. (1) for a QMA problem, with the additional promise that Ω does not have eigenvalues between a and b, $a - b > 1/\mathrm{poly}(n)$. Then the class #BQP consists of all problems of the form "compute the dimension of the space spanned by all eigenvectors with eigenvalues $\geq a$".

An equivalent definition for #BQP (cf. also [6, 7]) is the following: Consider a verifier Ω with the additional promise that there exist subspaces $\mathcal{A} \oplus \mathcal{R} = \mathbb{C}^{2^n}$ such that $\langle\psi|\Omega|\psi\rangle \geq a$ for all $|\psi\rangle \in \mathcal{A}$, and $\langle\psi|\Omega|\psi\rangle \leq b$ for all $|\psi\rangle \in \mathcal{R}$, where again $a - b > 1/\mathrm{poly}(n)$ – we can think of $\mathcal{A}$ and $\mathcal{R}$ as containing the good and bad witnesses, respectively. (Note that there will always be "mediocre" witnesses—the question is whether there exists a decomposition into a good and a bad witness space.) Then, #BQP consists of all problems of the form "compute dim $\mathcal{A}$". This number is well-defined, i.e., independent of the choice of $\mathcal{A}$ and $\mathcal{R}$, and moreover, one can easily show that it is equivalent to the definition above, cf. the Supplementary Material.

The gap promise we impose on the spectrum of Ω is not present in the definition of QMA (though similarly restricted versions of QMA were defined in [6, 7]). Nevertheless, this promise emerges naturally when considering the counting version: QMA captures the difficulty of determining the existence of an input state with acceptance probability above a, up to a "grace interval" $[b, a]$ in which mistakes are tolerated (i.e., if the largest eigenvalue of Ω is in $[b, a]$, the oracle can return either outcome). Correspondingly, #BQP captures the difficulty of counting the number of eigenvalues above a, where eivenvalues in the grace interval $[b, a]$ can be miscounted. The reason why we choose to define #BQP with a gap promise rather than with a grace interval is the same as for QMA, namely to have a unique outcome associated with any input.

Similarly, the idea of the Hamiltonian formulation of the problem which we will discuss below is to ask for the number of eigenstates in a certain energy interval, where states which are in some small $1/\mathrm{poly}(n)$ neighborhood of this interval may be miscounted; again, for reasons of rigor we choose to consider only Hamiltonians with no eigenstates in that interval. It should be noted, however, that all of the equivalence proofs we give equally apply if we choose to allow for miscounting of states in those grace intervals instead of requiring them to be empty, as the proofs do not make use of the gap promise itself, but rather show that all states outside those grace intervals are mapped (and thus counted) correctly. Thus, while the actual number returned by the grace interval formulation of the counting problems might change under those mappings due to different treatment of states in the grace interval, it will still be in the correct range.

The class #BQP inherits the important property from QMA of being stable under amplification, that is, the definition of #BQP is not sensitive to the choice of a and b. In particular, any $a - b > 1/\mathrm{poly}(n)$ can be amplified (by building a new poly-size Ω' from Ω) such that $a' = 1 - \exp(-\mathrm{poly}(n))$, $b' = \exp(-\mathrm{poly}(n))$, and keeping the eigenvalue gap between a' and b', by using a construction called *strong amplification*, cf. Ref. [8]; as shown there, strong amplification acts on all eigenvalues independently and thus also applies to #BQP. The crucial point is that strong amplification works without changing the proof itself, compared to weak amplification which takes multiple copies of the proof as an input. While this is fine for QMA, it does change the dimension of the accepting subspace in an unpredictable way and is thus not an option for the amplification of #BQP.

Complexity of computing the density of states.—Let us now show why the class #BQP is relevant for physical applications. In particular, we are going to show that computing the density of states of a local n-spin Hamiltonian $H = \sum_i H_i$ with few-body terms H_i, $\|H_i\| \leq 1$, up to accuracy $1/\mathrm{poly}(n)$, is a problem which is complete for #BQP, i.e., it is as hard as any problem in #BQP can be. The same holds true for the (a priori weaker) problem of counting the ground state degeneracy of a local Hamiltonian, given a $1/\mathrm{poly}(n)$ spectral gap above (note that Bravyi *et al.* [9] suggested this as a definition for a quantum counting class). We can impose additional restrictions on the interaction structure of our Hamilto-

nian, and as we will see, the hardness is preserved even for 2D lattices of qubits, or 1D systems.

The problem DOS (density of states) is defined as follows: Given a local Hamiltonian $H = \sum_i H_i$, compute the number of orthogonal eigenstates with eigenvalues in an interval $[E_1, E_2]$ with $E_2 - E_1 > 1/\mathrm{poly}(n)$, where we do not allow for eigenvalues within a small grace interval of width $\Delta = (E_2 - E_1)/\mathrm{poly}(n)$ centerd around E_1 and E_2; alternatively, we can allow for errorneous counts of eigenstates in that interval. Second, the problem #LH (sharp local Hamiltonian) corresponds to counting the number of ground states of a local Hamiltonian which has a spectral gap $\Delta = 1/\mathrm{poly}(n)$ above the ground state subspace, given we are told the ground state energy, and where we allow for a small splitting in the ground state energies; again, we can alternatively allow to miscount states in the grace interval.

Clearly, #LH is a special instance of DOS, i.e., solving #LH can be reduced to solving DOS. In order to show that DOS is contained in #BQP, we can use a phase estimation circuit [10] to estimate the energy of any given input $|\psi\rangle$ and only accept if its energy $\langle\psi|H|\psi\rangle$ is in the interval $[E_1, E_2]$; as the desired accuracy $\Delta = 1/\mathrm{poly}(n)$, this can be done efficiently. A detailed proof (using a more elementary circuit) is given in the Supplementary Material.

Let us now conversely show that #LH is a hard problem for #BQP, that is, any problem in #BQP can be reduced to counting the ground states of some gapped local Hamiltonian [22]. As in turn #LH can be reduced to DOS, which is contained in #BQP, this proves that both #LH and DOS are complete problems for #BQP, i.e., they capture the full difficulty of this class. To this end, we need to start from an arbitrary verifier circuit $U = U_T \cdots U_1$ and construct a Hamiltonian which has as many ground states as the circuit has accepting inputs (corresponding to the outcome $|1\rangle_1$ on the first qubit). Let $\mathcal{A}$ and $\mathcal{R}$ be the eigenspaces of Ω [Eq. (1)] with eigenvalues $\geq a = 1 - 2^{-\mathrm{poly}(n)}$ and $\leq b = 2^{-\mathrm{poly}(n)}$, respectively, and define $U[\mathcal{R}] := \{U|\psi\rangle_I|\mathbf{0}\rangle_A : |\psi\rangle_I \in \mathcal{R}\}$.

We will follow Kitaev's original construction for a Hamiltonian encoding a QMA verifier circuit [4, 5], which for any proof $|\psi\rangle_I \in \mathcal{A}$ has the "proof history" $|\Phi\rangle = \sum_{t=0}^{T} U_t \cdots U_1|\psi\rangle_I|\mathbf{0}\rangle_A|t\rangle_T$ as its ground state, where the third register is used as a clock. The Hamiltonian $H = H_{\mathrm{init}} + \sum_{t=1}^{T} H_{\mathrm{evol}}(t) + H_{\mathrm{final}}$ has three types of terms: $H_{\mathrm{init}} = \mathbb{1} \otimes (\mathbb{1} - |\mathbf{0}\rangle\langle\mathbf{0}|_A) \otimes |0\rangle\langle 0|_T$ makes sure the ancilla is initialized, $H_{\mathrm{evol}}(t) = -U_t \otimes |t\rangle\langle t-1|_T + \mathrm{h.c.}$ ensures proper evolution from $t-1$ to t, and $H_{\mathrm{final}} = \Pi_{U[\mathcal{R}]} \otimes |T\rangle\langle T|_T$ gives an energy penalty to states $|\Phi\rangle$ built from proofs $|\psi\rangle_I \in \mathcal{R}$. Note that our H_{final} differs from the usual choice $|0\rangle\langle 0|_1 \otimes \mathbb{1} \otimes |T\rangle\langle T|_T$ and is in fact non-local; as we show in the Supplementary Material, this does not significantly change the relevant spectral properties (in particular, we keep the $1/\mathrm{poly}(n)$ gap, and the ground state subspace is split at most exponentially). With this choice of H_{final}, H acts independently on the subspaces spanned by $\{U_t \cdots U_1|\psi\rangle_I|\boldsymbol{x}\rangle_A|t\rangle_T\}_{t=0,\ldots,T}$ for any $|\psi\rangle \in \mathcal{A}$ or $|\psi\rangle \in \mathcal{R}$, and $|\boldsymbol{x}\rangle_A$ the computational basis, and the restriction of H to any of these subspaces describes a random walk which is characterized by the choice of $|\psi\rangle_I$ and the number of 1's in $|\boldsymbol{x}\rangle_A$. These cases can be analyzed independently (see Supplementary Material), and it follows that H has a $\dim \mathcal{A}$-fold degenerate ground state space with a $1/\mathrm{poly}(n)$ gap above, proving #BQP-hardness of #LH.

This shows that #LH is #BQP-hard for a Hamiltonian which is a sum of $\log T$-local terms (i.e., each term acts on $\log T$ sites), as the clock register is of size $\log T$. In order to obtain a k-body Hamiltonian, Kitaev suggested to use a unary encoding of the clock (i.e., $|t\rangle_T$ is encoded as $|1\cdots 10\cdots 0\rangle$, with t 1's), so that each Hamiltonian term only acts on three qubits of the clock. However, this makes the Hilbert space of the clock too big, and terms need to be added to the Hamiltonian to penalize illegal clock configurations. These terms divide the Hilbert space into two parts, $\mathcal{H}_{\mathrm{legal}} \oplus \mathcal{H}_{\mathrm{illegal}}$. Here, $\mathcal{H}_{\mathrm{legal}}$ contains only legal clock states, whereas $\mathcal{H}_{\mathrm{illegal}}$ contains only configurations with illegal clock states [4, 5]. Since no Hamiltonian term couples these two subspaces, the Hamiltonian can be analyzed independently on the two subspaces. It turns out that its restriction to $\mathcal{H}_{\mathrm{illegal}}$ has an at least $1/\mathrm{poly}(n)$ higher energy, while on $\mathcal{H}_{\mathrm{legal}}$, the Hamiltonian is exactly the same as before. Thus, one finds that the new Hamiltonian still has the right number of ground states, and a $1/\mathrm{poly}(n)$ spectral gap. The very same argument applies in the case of 1D Hamiltonians, using the QMA-construction of Ref. [2]: Again, the Hamiltonian acts independently on a "legal" and an "illegal" subspace, where the latter has a polynomially larger energy, and the former reproduces the (low-energy) spectrum of the original Hamiltonian [11].

An alternative way to prove QMA-hardness on restricted interaction graphs is to use so-called *perturbation gadgets*, which yield the Hamiltonian of the Kitaev construction above from a perturbative expansion; in particular, this way one can show QMA-hardness of Pauli-type nearest-neighbor Hamiltonians on a 2D square lattice of qubits [1]. As shown in Ref. [12], such gadgets do in fact approximately preserve the whole low-energy part of the spectrum, and thus, our #BQP–hardness proof for #LH still applies to these classes of Hamiltonians. It should be noted, however, that since the eigenvalues are only preserved up to an error $1/\mathrm{poly}(n)$, the splitting of the ground state space will now be of order $1/\mathrm{poly}(n)$; however, it can still be chosen to be polynomially smaller than the spectral gap.

Quantum vs. classical counting complexity.—As we have seen, the quantum counting class #BQP exactly captures the difficulty of counting the degeneracy of ground states and computing the density of states of local quantum Hamiltonians. In the following, we will relate

#BQP to classical counting classes and prove that #BQP is equal to #P, counting the number of satisfying inputs to a boolean function [23]. In physical terms, this shows that counting the number of ground states or determining the density of states for a quantum Hamiltonian is not harder than either problem is for a classical Hamiltonian.

Clearly, #P is contained in #BQP, as the latter includes classical verifiers. It remains to be shown that any #BQP problem can be solved by computing a #P function. We start from a verifier operator Ω, Eq. (1), and wish to show that the dimension of its accepting subspace, i.e., the subspace $\mathcal{A}$ with eigenvalues $\geq a$, can be computed by counting satisfying inputs to some efficiently computable boolean function. Using amplification, we can ensure that $|\dim \mathcal{A} - \operatorname{tr} \Omega| \leq \frac{1}{4}$, i.e., we need to compute $\operatorname{tr} \Omega$ to accuracy $\frac{1}{4}$. This can be done using a "path integral" method, which has been used previously to show containments of quantum classes in the classical classes PP and #P (see e.g. [13]). We rewrite $\operatorname{tr} \Omega = \sum_\zeta f(\zeta)$ as a sum over products of transition probabilities along a path $\zeta \equiv (i_0, \ldots, i_N, j_1, \ldots, j_N)$, where

$$f(\zeta) = \langle i_0|_I \langle \mathbf{0}|_A U_1^\dagger |j_1\rangle\langle j_1| U_1^\dagger \cdots U_T^\dagger |j_T\rangle \times \langle i_T| \left[|0\rangle\langle 0|_1 \otimes \mathbb{1} \right] |i_T\rangle\langle i_T| U_T \cdots U_1 |i_0\rangle_I |\mathbf{0}\rangle_A \quad (2)$$

(cf. Fig. 1). Since any such sum over an efficiently computable $f(\zeta)$ can be determined by counting the satisfying inputs to some boolean formula (see the Supplementary Material for details), it follows that Ω can be computed using a single query to a black box solving #P problems.

Summary and discussion.—In this work, we considered two problems: Computing the density of states and computing the ground state degeneracy of a local Hamiltonian of a spin system. In order to capture the computational difficulty of these problems we introduced the quantum complexity class #BQP, the counting version of the class QMA. We proved that this complexity class exactly captures the difficulty of our two problems, even when restricting to local Hamiltonians on 2D lattices of qubits or to 1D chains, since all these problems are complete problems for the class #BQP [24].

We have further shown that #BQP is no harder than its classical counterpart #P. In particular this implies that computing the density of states is no harder for quantum Hamiltonians than it is for classical ones. While this quantum-classical equivalence might seem surprising at the Hamiltonian level, it should be noted that the classes #P and PP quite often form natural "upper bounds" for many quantum *and* classical problems.

What about the problem of computing the density of states for fermionic systems, such as many-electron systems? On the one hand, this problem will be still in #BQP and thus #P, since any local fermionic Hamiltonian can be mapped via the Jordan-Wigner transform to a (non-local) Hamiltonian on a spin system, whose energy can still be estimated efficiently by a quantum circuit [14]. On the other hand, hardness of the problem for #BQP can be shown e.g. by using the #BQP-hardness of #LH, and encoding each spin using one fermion in two modes, similar to [14]. Thus, computing the density of states for fermionic systems is a #BQP-complete problem as well.

Acknowledgments.—We thank S. Aaronson, D. Gottesman, D. Gross, Y. Shi, M. van den Nest, J. Watrous and S. Zhang for helpful discussions and comments. Research at PI is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. NS is supported by the Gordon and Betty Moore Foundation through Caltech's Center for the Physics of Information and NSF Grant No. PHY-0803371.

Note added.—After completion of this work, we learned that Shi and Zhang [15] have independently defined #BQP and shown its relation to #P using the same technique.

[1] R. Oliveira and B. M. Terhal, Quant. Inf. Comput. **8**, 900 (2009), quant-ph/0504050.
[2] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe, Commun. Math. Phys. **287**, 41 (2009), arXiv:0705.4077.
[3] N. Schuch and F. Verstraete, Nature Physics **5**, 732 (2009), arXiv:0712.0483.
[4] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and quantum computation* (American Mathematical Society, Providence, Rhode Island, 2002).
[5] D. Aharonov and T. Naveh, (2002), quant-ph/0210077.
[6] D. Aharonov, M. Ben-Or, F. G. Brandao, and O. Sattath, (2008), arXiv:0810.4840.
[7] R. Jain, I. Kerenidis, M. Santha, and S. Zhang, (2009), arXiv:0906.4425.
[8] C. Marriott and J. Watrous, Comput. Complex. **14**, 122 (2005), cs/0506068.
[9] S. Bravyi, C. Moore, and A. Russell, (2009), arXiv:0907.1297.
[10] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Proc. R. Soc. Lond. **454**, 339 (1998), quant-ph/9708016.
[11] N. Schuch, I. Cirac, and F. Verstraete, Phys. Rev. Lett. **100**, 250501 (2008), arXiv:0802.3351.
[12] J. Kempe, A. Kitaev, and O. Regev, SIAM Journal of Computing **35**, 1070 (2006), quant-ph/0406180.
[13] L. Adleman, J. DeMarrais, and M.-D. Huang, SIAM J. Comput. **26**, 1524 (1997).
[14] Y.-K. Liu, M. Christandl, and F. Verstraete, Phys. Rev. Lett. **98**, 110503 (2007), quant-ph/0609125.
[15] Y. Shi and S. Zhang, private communication; notes available at `http://www.cse.cuhk.edu.hk/~syzhang/papers/SharpBQP.pdf`.
[16] R. Bhatia, *Matrix Analysis* (Springer, 1997).
[17] C. Jordan, Bulletin de la Société Mathématique de France **3**, 103 (1875).
[18] P. Halmos, Trans. Amer. Math. Soc. **144**, 381 (1969).
[19] E. Bernstein and U. Vazirani, Siam J. of Comp. **26**, 1411 (1997), quant-ph/9701001.
[20] F.G.S.L. Brandao, PhD thesis (2008), arXiv:0810.0026.

[21] T. J. Osborne, (2006), cond-mat/0605194.
[22] Note that the connection between QMA with a unique "good witness", such as in our #BQP definition, and local Hamiltonians with unique ground state and a $1/\mathrm{poly}(n)$ gap has been shown in [6].
[23] Formally speaking, the reduction from #BQP to #P is *weakly parsimonious*, i.e., for any function $f \in \#\mathsf{BQP}$ there exist polynomial-time computable functions α and β, and a function $g \in \#\mathsf{P}$, such that $f = \alpha \circ g \circ \beta$. This differs from Karp reductions where no postprocessing is allowed, $\alpha = \mathrm{Id}$, but it still only requires a single call to a #P oracle.
[24] Note that computing the density of states to multiplicative accuracy is less difficult, see Refs. [20, 21].

SUPPLEMENTARY MATERIAL

Quantum Complexity Classes

We start with the relevant definitions. Let x be a binary string. Then, we denote by the verifier $V \equiv V_x$ a quantum circuit of length $T = \mathrm{poly}(|x|)$, $U = U_T \cdots U_1$ (with local gates U_t) acting on $m = \mathrm{poly}(|x|)$ qubits, which is generated uniformly from x. The verifier takes an $n = \mathrm{poly}(|x|)$ qubit quantum state $|\psi\rangle_I$ as an input (we will express everything in terms of n instead of $|x|$ in the following), together with $m - n$ initialized ancillas, $|\mathbf{0}\rangle_A \equiv |0 \cdots 0\rangle_A$, applies U, and finally measures the first qubit in the $\{|0\rangle_1, |1\rangle_1\}$ basis to return 1 ("proof accepted") or 0 ("proof rejected"). The acceptance probability for a proof $|\psi\rangle$ is then given by $p_{\mathrm{acc}}(V(\psi)) := \langle\psi|\Omega|\psi\rangle$, with

$$\Omega = (\mathbb{1}_I \otimes \langle\mathbf{0}|_A) U^\dagger (|1\rangle\langle 1|_1 \otimes \mathbb{1}) U (\mathbb{1}_I \otimes |\mathbf{0}\rangle_A)\,. \tag{3}$$

Definition 1 (QMA)**.** *Let $0 \le a, b \le 1$ s.th. $a - b > \frac{1}{p(n)}$ for some polynomial $p(n)$. A language L is in $\mathsf{QMA}(a,b)$ if there exists a verifier s.th.*

$$\begin{aligned} x \in L &\Rightarrow \lambda_{\max}(\Omega) > a \\ x \notin L &\Rightarrow \lambda_{\max}(\Omega) < b\,. \end{aligned}$$

Definition 2 (#BQP)**.** *Let $0 \le a, b \le 1$ s.th. $a - b > 1/\mathrm{poly}(n)$, and let Ω be a verifier map with no eigenvalues between a and b. Then, the class $\#\mathsf{BQP}(a,b)$ consists of all problems of the form "compute the dimension of the space spanned by all eigenvectors of Ω with eigenvalues $\ge a$" .*

An alternative definition for #BQP is the following:

Definition 3 (#BQP, alternate definition)**.** *Consider a verifier Ω with the property that there exist subspaces $\mathcal{A} \oplus \mathcal{R} = \mathbb{C}^N$ ($N = 2^n$) such that $\langle\psi|\Omega|\psi\rangle \ge a$ for all $|\psi\rangle \in \mathcal{A}$, and $\langle\psi|\Omega|\psi\rangle \le b$ for all $|\psi\rangle \in \mathcal{R}$, where again $a - b > 1/\mathrm{poly}(n)$. Then $\#\mathsf{BQP}(a,b)$ consists of all problems of the form "compute* $\dim \mathcal{A}$*".*

Note that $\dim \mathcal{A}$ is well-defined: Consider two decompositions $\mathbb{C}^N = \mathcal{A} \oplus \mathcal{R}$ and $\mathbb{C}^N = \mathcal{A}' \oplus \mathcal{R}'$. Without loss of generality, if we assume $\dim \mathcal{A} > \dim \mathcal{A}'$, it follows $\dim \mathcal{A} + \dim \mathcal{R}' > N$, and thus there exists a non-trivial $|\mu\rangle \in \mathcal{A} \cap \mathcal{R}'$, which contradicts the definition.

Theorem 4. *Definition 2 and Definition 3 are equivalent.*

Proof. To show that Definition 2 implies Definition 3, let $\mathcal{A}$ be spanned by the eigenvectors with eigenvalues $\ge a$. To show the converse, we use the minimax principle for eigenvalues [16], which states that the kth largest eigenvalue λ_k of a Hermitian operator Ω in an N-dimensional Hilbert space can be obtained from either of the equivalent optimizations

$$\lambda_k(\Omega) = \max_{\mathcal{M}_k} \min_{|x\rangle \in \mathcal{M}_k} \langle x|\Omega|x\rangle \tag{4}$$

$$= \min_{\mathcal{M}_{N-k+1}} \max_{|x\rangle \in \mathcal{M}_{N-k+1}} \langle x|\Omega|x\rangle\,, \tag{5}$$

where $\mathcal{M}_k$ is a subspace of dimension k, and $|x\rangle$ is a unit vector. Now notice that Def. 3 implies that

$$\min_{|x\rangle \in \mathcal{A}} \langle x|\Omega|x\rangle \ge a \quad \text{and} \quad \max_{|x\rangle \in \mathcal{R}} \langle x|\Omega|x\rangle \le b\,. \tag{6}$$

Next, consider the minimax theorem for $k = \dim \mathcal{A}$. From Eq. (4) we have

$$\begin{aligned} \lambda_{\dim \mathcal{A}} &= \max_{\mathcal{M}_{\dim \mathcal{A}}} \min_{|x\rangle \in \mathcal{M}_{\dim \mathcal{A}}} \langle x|\Omega|x\rangle \\ &\ge \min_{|x\rangle \in \mathcal{A}} \langle x|\Omega|x\rangle \\ &\ge a\,. \end{aligned}$$

Now consider the case that $k = \dim \mathcal{A} + 1$. From Eq. (5), using the fact that $N - (\dim \mathcal{A} + 1) + 1 = \dim \mathcal{R}$, we have

$$\begin{aligned} \lambda_{\dim \mathcal{A}+1} &= \min_{\mathcal{M}_{\dim \mathcal{R}}} \max_{|x\rangle \in \mathcal{M}_{\dim \mathcal{R}}} \langle x|\Omega|x\rangle \\ &\le \max_{|x\rangle \in \mathcal{R}} \langle x|\Omega|x\rangle \\ &\le b\,. \end{aligned}$$

Thus we have

$$\lambda_{\dim \mathcal{A}} \ge a > b \ge \lambda_{\dim \mathcal{A}+1}\,, \tag{7}$$

since $a - b \geq 1/\mathrm{poly}(n)$. It follows that $\lambda_{\dim \mathcal{A}}$ is the smallest eigenvalue of Ω which is still larger than a, and therefore the span of the first $\dim \mathcal{A}$ eigenvectors of Ω is equal to the span of all eigenvectors with eigenvalue $\geq a$. The equivalence follows. □

The class $\#\mathsf{BQP}(a, b)$ inherits the useful property of strong error reduction from QMA: the thresholds (a, b) can be amplified to $(1 - 2^{-r}, 2^{-r})$, $r = \mathrm{poly}(n)$ while keeping the size of the proof:

Theorem 5. *Let $\#\mathsf{BQP}_n(a, b)$ denote $\#\mathsf{BQP}$ with an n qubit witness. Then $\#\mathsf{BQP}_n(a, b) \subset \#\mathsf{BQP}_n(1 - 2^{-r}, 2^{-r})$ for every $r \in \mathrm{poly}(n)$.*

This follows directly from the strong amplification procedure presented in [8], which describes a procedure to amplify any verifier map Ω such that any eigenvalue above a (below b) is shifted above $1-2^{-r}$ (below 2^{-r}) at an overhead polynomial in r.

Using Thm. 5, we will write $\#\mathsf{BQP}$ instead of $\#\mathsf{BQP}(a, b)$ from now on, where $a - b > \mathrm{poly}(n)$, and a, b can be exponentially close to 1 and 0, respectively.

The Complexity of the Density of States

We now use the class $\#\mathsf{BQP}$ to characterize the complexity of the density of states problem and the problem of counting the number of ground states of a local Hamiltonian. We start by defining these problems, as well as the notion of local Hamiltonian, and then show that both problems are $\#\mathsf{BQP}$-complete.

Definition 6 (k-local Hamiltonian)**.** *Given a set of $\mathrm{poly}(n)$ quantum spins each with dimension bounded by a constant, a Hamiltonian H for the system is said to be k-local if $H = \sum_i H_i$ is a sum of at most $\mathrm{poly}(n)$ Hermitian operators H_i, $\|H_i\| \leq 1$, each of which acts nontrivially on at most k spins.*

Note that k-local does not imply any geometric locality, only that each spin interacts with at most $k-1$ other spins for any given interaction term. However, we restrict ourselves to $k = O\big(\log(n)\big)$ so that each H_i can be specified by an efficient classical description.

Definition 7 (DOS)**.** *Let $E_2 - E_1 > 1/\mathrm{poly}(n)$, $\Delta = (E_2 - E_1)/\mathrm{poly}(n)$, and let $H = \sum_i H_i$ be a k-local Hamiltonian such that H has no eigenvalues in the intervals $[E_1 - \frac{\Delta}{2}, E_1 + \frac{\Delta}{2}]$ and $[E_2 - \frac{\Delta}{2}, E_2 + \frac{\Delta}{2}]$. Then, the problem* DOS *(density of states) is to compute the number of orthogonal eigenstates with eigenvalues in the interval $[E_1, E_2]$.*

Definition 8 (#LH)**.** *Let $E_0 \leq E_1 < E_2$, with $E_2 - E_1 > 1/\mathrm{poly}(n)$, and let $H = \sum_i H_i$ be a k-local Hamiltonian s.th. $H \geq E_0$, and H has no eigenvalues between E_1 and E_2. Then, the problem $\#\mathrm{LH} \equiv \#\mathrm{LH}(E_1 - E_0)$ (sharp local Hamiltonian) is to compute the dimension of the eigenspace with eigenvalues $\leq E_1$.*

Note that #LH depends on the "energy splitting" $E_1 - E_0$ of the low-energy subspace. In particular, for $E_1 - E_0 = 0$, $\#\mathrm{LH}(0)$ corresponds to computing the degeneracy of the ground state subspace. As we will see in what follows, the class $\#\mathrm{LH}(\sigma)$ is the same for any splitting $\exp(-\mathrm{poly}(n)) \leq \sigma \leq 1/\mathrm{poly}(n)$.

We now show that #LH and DOS are both $\#\mathsf{BQP}$-complete. We do so by giving reductions from $\#\mathrm{LH}(1/\mathrm{poly}(n))$ to DOS, from DOS to $\#\mathsf{BQP}$, and from $\#\mathsf{BQP}$ to $\#\mathrm{LH}\big(\exp(-\mathrm{poly}(n))\big)$; this will at the same time prove the claimed independence of $\#\mathrm{LH}(\sigma)$ of the splitting σ.

Theorem 9. $\#\mathrm{LH}(1/\mathrm{poly}(n))$ *reduces to* DOS.

Proof. If we denote the parameters of the #LH problem by $\tilde{E}_0, \tilde{E}_1, \tilde{E}_2$, then we can simply relate them to the parameters E_1, E_2, Δ of a DOS problem by $\Delta = \tilde{E}_2 - \tilde{E}_1$, $E_1 = \tilde{E}_0 - \frac{1}{2}\Delta$ and $E_2 = \tilde{E}_1 + \frac{1}{2}\Delta$, and the result follows directly. □

Theorem 10. DOS *is contained in* $\#\mathsf{BQP}$.

Proof. We start with a k-local Hamiltonian H as in Def. 7. Now define a new Hamiltonian

$$H' := \nu(H^2 - (E_1 + E_2)H + E_1 E_2) \ .$$

H' is a $2k$-local Hamiltonian; here, $\nu = 1/\mathrm{poly}(n)$ is chosen such that each term in H' is subnormalized. Any eigenvalue of H in the interval $[E_1 + \frac{\Delta}{2}; E_2 - \frac{\Delta}{2}]$ translates into an eigenvalue of H' which is below

$$A := -\nu\tfrac{\Delta}{2}(E_2 - E_1 - \tfrac{\Delta}{2}) \leq -1/\mathrm{poly}(n) \ ,$$

whereas any eigenvalue outside $[E_1 - \frac{\Delta}{2}; E_2 + \frac{\Delta}{2}]$ translates into an eigenvalues of H' above

$$B := \nu\tfrac{\Delta}{2}(E_2 - E_1 + \tfrac{\Delta}{2}) \geq 1/\mathrm{poly}(n) \ .$$

The original DOS problem now translates into counting the number of eigenstates of H' with negative energy, given a spectral gap in a $1/\mathrm{poly}(n)$ sized interval $[A; B]$ around zero. We now use the circuit which was used in [5] to prove that log-local Hamiltonian is in QMA; it accepts any input state $|\psi\rangle$ with probability

$$p(|\psi\rangle) = \frac{1}{2} - \frac{\langle\psi|H'|\psi\rangle}{2m} \ ,$$

where $m = \mathrm{poly}(n)$ is the number of terms in H'. [The idea is to randomly pick one term H'_i in the Hamiltonian and toss a coin with probability $(1 - \langle\psi|H'_i|\psi\rangle)/2$.] Computing the answer to the original DOS problem then translates to counting the number of states with

acceptance probability $\geq a = \frac{1}{2} + \frac{A}{2m}$, given that there are no eigenstates between a and $b = \frac{1}{2} + \frac{B}{2m}$, and since $a - b = (A - B)/2m \geq 1/\mathrm{poly}(n)$, this shows that this number can be computed in #BQP. □

Theorem 11. *$\#\mathrm{LH}\big(\exp(-\mathrm{poly}(n))\big)$ is #BQP-hard.*

Proof. To show #BQP-hardness of #LH, we need to start with an arbitrary QMA verifier circuit $U = U_T \dots U_1$ and construct a Hamiltonian with as many ground states as the circuit has accepting inputs. By amplification, we can assume that the acceptance and rejection thresholds for the verifier are $a = 1 - \epsilon$ and $b = \epsilon$, where we can choose $\epsilon = O(\exp(-p(n)))$ for any polynomial $p(n)$. As before, let $\mathcal{A}$ and $\mathcal{R}$ be the eigenspaces of Ω with eigenvalues $\geq a$ and $\leq b$, respectively. Define

$$U[\mathcal{R}] := \{U|\psi\rangle_I|\mathbf{0}\rangle_A : |\psi\rangle_I \in \mathcal{R}\} \tag{8}$$

and denote the projector onto this space by $\Pi_{U[\mathcal{R}]}$. Notice that for any state $|\chi\rangle \in U[\mathcal{R}]$, due to our rejection threshold $b = \epsilon$, we have

$$\langle\chi|\big(|1\rangle\langle 1|_1 \otimes \mathbb{1}\big)|\chi\rangle \leq \epsilon\,. \tag{9}$$

We now follow Kitaev's original construction to encode a QMA verifier circuit into a Hamiltonian which has a "proof history" as its ground state for any proof $|\phi\rangle_I \in \mathcal{A}$ [4, 5]. That is, the ground states of the Hamiltonian are given by

$$|\Phi\rangle = \sum_{t=0}^{T} U_t \dots U_1|\phi\rangle_I|\mathbf{0}\rangle_A|t\rangle_T \tag{10}$$

where the third register is used as a "clock". The Hamiltonian has the form

$$H = H_{\mathrm{init}} + \sum_{t=1}^{T} H_{\mathrm{evol}}(t) + H_{\mathrm{final}} \tag{11}$$

where

- $H_{\mathrm{init}} = \mathbb{1}_I \otimes (\mathbb{1} - |\mathbf{0}\rangle\langle\mathbf{0}|_A) \otimes |0\rangle\langle 0|_T$ checks that the ancilla is property initialized, penalizing states without properly initialized ancillas;
- $H_{\mathrm{evol}}(t) = -\frac{1}{2}U_t \otimes |t\rangle\langle t-1|_T - \frac{1}{2}U_t^\dagger \otimes |t-1\rangle\langle t|_T$ $+\frac{1}{2}\mathbb{1} \otimes |t\rangle\langle t|_T + \frac{1}{2}\mathbb{1} \otimes |t-1\rangle\langle t-1|_T$ checks that the propagation from time $t-1$ to t is correct, penalizing states with erroneous propagation;
- $H_{\mathrm{final}} = \Pi_{U[\mathcal{R}]} \otimes |T\rangle\langle T|_T$ causes each state $|\phi\rangle$ built from an input $|\psi\rangle_I \in \mathcal{R}$ (but which nonetheless has a correctly initialized ancilla) to receive an energy penalty.

As we show in Lemma 12, the total Hamiltonian H has a spectral gap $1/\mathrm{poly}(n)$ above the $\dim \mathcal{A}$–dimensional ground state subspace. However, H_{final} is not a local Hamiltonian, but as we argue in the following, it can be replaced by the usual term $H^{\mathrm{std}}_{\mathrm{final}} = |0\rangle\langle 0|_1 \otimes \mathbb{1} \otimes |T\rangle\langle T|_T$ while keeping the ground space dimension (up to small splitting in energies) and the $1/\mathrm{poly}(n)$ spectral gap. As we prove in Lemma 13,

$$H^{\mathrm{std}}_{\mathrm{final}} \geq H_{\mathrm{final}} - \sqrt{\epsilon}\mathbb{1}\,. \tag{12}$$

Thus, replacing H_{final} by $H^{\mathrm{std}}_{\mathrm{final}}$ will decrease the energy of any excited state by at most $\sqrt{\epsilon} = O(\exp(-p(n)/2))$. (The energy of the ground states is already minimal and cannot decrease.) On the other hand, the energy of any proper proof history for H cannot increase by more than

$$\langle\chi|H^{\mathrm{std}}_{\mathrm{final}}|\chi\rangle = \langle\chi||0\rangle\langle 0|_1 \otimes \mathbb{1}|\chi\rangle \leq \epsilon = O(\exp(-p(n))\,,$$

due to our choice of acceptance threshold, i.e., the low energy subspace has dimension $\dim \mathcal{A}$. We thus obtain a Hamiltonian with a $\dim \mathcal{A}$ dimensional ground state subspace with energy splitting $\leq \epsilon = \exp(-p(n))$ and a $1/\mathrm{poly}(n)$ spectral gap above. □

The following two lemmas are used in the preceding proof of Theorem 11.

Lemma 12. *H has a spectral gap of size $1/\mathrm{poly}(n)$.*

Proof. Our proof follows closely the discussion in Ref. [4], cf. also [11]. We can block diagonalize H by the (conjugate) action of the following unitary operator,

$$W = \sum_{j=0}^{T} U_j \cdots U_1 \otimes |j\rangle\langle j|_T\,, \tag{13}$$

which maps $H \to H' = W^\dagger H W$. As this has no effect on the spectrum, we can work with the simpler H' from now on.

Let us explicitly write the effect of conjugation by W on the terms of H. The first term is unaffected, $H'_{\mathrm{init}} = H_{\mathrm{init}}$. The final term becomes $H'_{\mathrm{final}} = \Pi_{\mathcal{R}} \otimes |\mathbf{0}\rangle\langle\mathbf{0}|_A \otimes |T\rangle\langle T|_T$, where $\Pi_{\mathcal{R}}$ is simply the projector onto the space $\mathcal{R}$, and the ancillas are in the correct initial state.

We can now conjugate each of the terms in $H_{\mathrm{evol}}(t)$ separately. For example, the first term gives

$$W^\dagger(U_t \otimes |t\rangle\langle t-1|)W = \mathbb{1} \otimes |t\rangle\langle t-1|\,. \tag{14}$$

The other terms are exactly analogous, and we find that

$$H'_{\mathrm{evol}}(t) = \mathbb{1} \otimes \tfrac{1}{2}\big[|t-1\rangle\langle t-1| + |t\rangle\langle t| - |t\rangle\langle t-1| - |t-1\rangle\langle t|\big]\,.$$

The total evolution Hamiltonian is then block-diagonal a matrix which looks like a hopping Hamiltonian in the clock register,

$$\sum_t H'_{\mathrm{evol}}(t) = \mathbb{1} \otimes E\,, \tag{15}$$

where the $(T+1)$-by-$(T+1)$–dimensional tri-diagonal matrix E is given by

$$E = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & & & \\ -\frac{1}{2} & 1 & -\frac{1}{2} & & \\ & -\frac{1}{2} & \ddots & & \\ & & & 1 & -\frac{1}{2} \\ & & & -\frac{1}{2} & \frac{1}{2} \end{pmatrix} . \tag{16}$$

To discuss the spectrum of H', let

$$\mathcal{S}_1 = \mathcal{A} \otimes |0\rangle\langle 0|_A \otimes \mathbb{C}^{T+1} ,$$
$$\mathcal{S}_2 = \mathcal{R} \otimes |0\rangle\langle 0|_A \otimes \mathbb{C}^{T+1} ,$$

and $\mathcal{S}_3$ the orthogonal complement of $\mathcal{S}_1 \oplus \mathcal{S}_2$; i.e., $\mathcal{S}_1$ corresponds to evolutions starting from good proofs, $\mathcal{S}_2$ to those starting from wrong proofs, and $\mathcal{S}_3$ to evolutions with wrongly initialized ancillas. Note that $H' = H|_{\mathcal{S}_1} \oplus H|_{\mathcal{S}_2} \oplus H|_{\mathcal{S}_3}$, thus, we can analyze the spectrum for $H|_{\mathcal{S}_p}$ separately. Note that the restriction to $\mathcal{S}_p$ does not affect $H'_{\text{evol}}(t)$. Since we expect the ground state subspace to occur on $\mathcal{S}_1$, we need to compute ground state energy and gap of $H'|_{\mathcal{S}_1}$, as well as lower bound the ground state energies of $H'|_{\mathcal{S}_2}$ and $H'_{\mathcal{S}_3}$.

First, $H'_{\text{init}}|_{\mathcal{S}} = H'_{\text{final}}|_{\mathcal{S}} = 0$, i.e., the spectrum of $H'|_{\mathcal{S}}$ equals the spectrum of E, which can be straightforwardly determined to be $1 - \cos\omega_n$, with $\omega_n = n\pi/(T+1)$, and eigenfunctions $(\cos\frac{1}{2}\omega_n, \cos\frac{3}{2}\omega_n, \dots)$; the ground state degeneracy is indeed $\dim\mathcal{A}$ as desired.

On the other hand, $H'_{\text{final}}|_{\mathcal{S}_2} = 1$, and $H'_{\text{init}}|_{\mathcal{S}_3} \geq 1$, i.e., in both cases the ground state energy of $H'|_{\mathcal{S}_p}$ is lower bounded by the ground state energy of

$$E' = \begin{pmatrix} \frac{1}{2} & -\frac{1}{2} & & & \\ -\frac{1}{2} & 1 & -\frac{1}{2} & & \\ & -\frac{1}{2} & \ddots & & \\ & & & 1 & -\frac{1}{2} \\ & & & -\frac{1}{2} & \frac{3}{2} \end{pmatrix} ,$$

which has eigenvalues $1-\cos\vartheta_n$, with $\vartheta_n = (n+\frac{1}{2})\pi/(T+\frac{3}{2})$, and eigenfunctions $(\cos\frac{1}{2}\vartheta_n, \cos\frac{3}{2}\vartheta_n, \dots)$.

It follows that H' (and thus H) has a ground state energy of $1 - \cos\omega_0 = 0$, and a gap $1 - \cos\frac{\pi}{2T+3} = O(1/T^2) = O(1/\text{poly}(n))$ above. □

It remains to prove Eq. (12), which follows from the following Lemma by choosing $P = |0\rangle\langle 0|_1 \otimes \mathbb{1}$, $Q = \Pi_{U[\mathcal{R}]}$, and using Eq. (9).

Lemma 13. *Let P and Q be projectors such that $\|Q(\mathbb{1}-P)Q\|_\infty \leq \epsilon$. Then*

$$P - Q \geq -\sqrt{\epsilon}\mathbb{1} . \tag{17}$$

Proof. We begin by recalling the result due to Jordan [17] (see Ref. [18] for a more modern treatment) for the simultaneous canonical form of two projectors. In the subspace where P and Q commute, both operators are diagonal in a common basis and the spectrum is either $(0,0), (0,1), (1,0)$, or $(1,1)$, and direct sums of those terms. In the subspace where they don't commute, the problem decomposes into a direct sum of two-by-two blocks given by

$$P_j = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad , \quad Q_j = \begin{pmatrix} c^2 & cs \\ cs & s^2 \end{pmatrix} , \tag{18}$$

where $s = \sin(\theta_j)$ for some angle θ_j, $c^2 + s^2 = 1$, and the subscript j just labels a generic block.

In fact, this two-by-two block form is completely general if we allow embedding our projectors into a larger space while preserving their rank. The rank-preserving condition guarantees that our bound is unchanged, since we are only appending blocks of zeros, and so we will consider this two-by-two form without loss of generality.

The constraint that $\|Q(\mathbb{1} - P)Q\|_\infty \leq \epsilon$ implies constraints on the values that $\sin(\theta_j)$ can take. In particular, we can directly compute this operator norm in each block separately, and we find that for all j

$$\|Q_j(\mathbb{1} - P_j)Q_j\|_\infty = \sin^2(\theta_j) \leq \epsilon . \tag{19}$$

We can also directly compute in each block that

$$P_j - Q_j = \begin{pmatrix} 1-c^2 & 1-cs \\ -cs & -s^2 \end{pmatrix} . \tag{20}$$

The spectrum of this operator is easily computed to be $\pm|\sin(\theta_j)|$. Thus, the least eigenvalue of $P - Q$ is bounded from below by $-\sqrt{\epsilon}$. □

Quantum vs. Classical Counting Complexity

We finally want to relate #BQP to the classical counting class #P. It is clear that any #P problem can be solved in #BQP by using a classical circuit. We will now show that conversely, #BQP can be reduced to #P under *weakly parsimonious reductions*: That is, for any function $f \in$ #BQP there exist polynomial-time computable functions α and β, and a function $g \in$ #P, such that $f = \alpha \circ g \circ \beta$. This differs from Karp reductions where no postprocessing is allowed, $\alpha = \text{Id}$, but still only requires a single call to a #P oracle, in contrast to Turing reductions.

Theorem 14. *There exists a weakly parsimonious reduction from #BQP to #P.*

Proof. We start from a verifier operator Ω for a #BQP problem. First, we use strong error reduction to let $a =$

$1 - 2^{-(n-2)}$ and $b = 2^{-(n+2)}$. It follows that

$$|\dim \mathcal{A} - \operatorname{tr} \Omega| \leq 2^n 2^{-(n+2)} = \frac{1}{4} \tag{21}$$

and thus we need to compute $\operatorname{tr} \Omega$ to accuracy $1/4$. This can be done using the "path integral" method previously used to show containments of quantum classes inside PP and $\#\mathsf{P}$ [13]. We rewrite $\operatorname{tr} \Omega = \sum_\zeta f(\zeta)$, where the sum is over products of transition probabilities along a path, which we label

$$\zeta \equiv (i_0, \ldots, i_N, j_1, \ldots, j_N)\,, \tag{22}$$

so that

$$\begin{aligned} f(\zeta) =& \langle i_0|_I \langle \mathbf{0}|_A U_1^\dagger |j_1\rangle\langle j_1| U_1^\dagger \cdots U_T^\dagger |j_T\rangle \times \\ & \langle i_T| \big[|0\rangle\langle 0|_1 \otimes \mathbb{1} \big] |i_T\rangle\langle i_T| U_T \cdots U_1 |i_0\rangle_I |\mathbf{0}\rangle_A\,. \end{aligned} \tag{23}$$

(cf. Fig. 1 in the main manuscript for an illustration).

Since any quantum circuit can be recast in terms of real gates at the cost of doubling the number of qubits [19], we can simplify the proof by assuming $f(\zeta)$ to be real. To achieve the desired accuracy it is sufficient to approximate f up to $|\zeta| + 2$ digits, where $|\zeta| = \operatorname{poly}(n)$ is the number of bits in ζ. Now define

$$g(\zeta) := \operatorname{round}\big[2^{|\zeta|+2}(f(\zeta) + 1)\big] \tag{24}$$

and note that $g(\zeta)$ is a positive and integer-valued function satisfying

$$\Big| \big[2^{-|\zeta|-2} \sum_\zeta g(\zeta) - 1\big] - \sum_\zeta f(\zeta) \Big| \leq \tfrac{1}{4}\,. \tag{25}$$

Finally, by defining a boolean indicator function,

$$h(\zeta, \xi) = \begin{cases} 1 & \text{if } 0 \leq \xi < g(\zeta) \\ 0 & \text{otherwise} \end{cases}$$

we can write $g(\zeta) = \sum_{\xi \geq 0} h(\zeta, \xi)$, and thus,

$$\sum_\zeta g(\zeta) = \sum_{\xi,\zeta} h(\zeta, \xi)\ .$$

This shows that $\operatorname{tr} \Omega$ can be approximated to accuracy $\frac{1}{4}$, and thus $\dim \mathcal{A}$ can be determined by counting the number of satisfying assignments of a single boolean function $h(\zeta, \xi)$ that can be efficiently constructed from Ω, i.e., by a single query to a black box solving $\#\mathsf{P}$ problems, together with the efficient postprocessing described by Eq. (25). □